LA-UR-17-24755
Approved for public release; distribution is unlimited.

Title: Radiographic Inference Based on a Model of V1 Simple Cells Implemented on the D-Wave 2X Quantum Annealing Computer

Author(s): Nguyen, Nga Thi Thuy
Kenyon, Garrett

Intended for: IEEE International Conference on Rebooting Computing (ICRC), 2017-11-07/2017-11-09 (Washington, District Of Columbia, United States)

Issued: 2018-12-11 (rev.1)

Radiographic Inference Based on a Model of V1 Simple Cells Implemented on the D-Wave 2X Quantum Annealing Computer

Nga T. T. Nguyen
Information Sciences
Los Alamos National Laboratory
Los Alamos, New Mexico 87545, USA

Garrett T. Kenyon
Information Sciences
Los Alamos National Laboratory &
New Mexico Consortium
Los Alamos, New Mexico 87545, USA
gkenyon@lanl.gov

***Abstract*—Just as the brain must infer 3D structure from 2D retinal images, radiologists are tasked with inferring 3D densities from 2D X-rays. Computer simulations suggest that V1 simple cells use lateral inhibition to generate sparse representations that are selective for 3D depth when presented with 2D stereo images and video. Analogously, we cast radiographic inference as a sparse coding problem employing lateral inhibition between binary neurons, resulting in a quadratic unconstrained binary optimization (QUBO) problem suitable for implementation on a quantum annealing D-Wave 2X (1152-qubit) computer. We generated synthetic radiographs by performing discrete Abel transforms on mathematically-defined objects possessing axial (cylindrical) symmetry and whose radially density profile was given by the sum of a randomly-chosen, sparse set of (nearly binary) Fourier components. We used embedding tools to map the above QUBO problem, which involved dense connections between up to 47 Fourier coefficients, onto the very sparsely connected D-Wave chimera. Using quantum inference, we were able to reconstruct reasonably accurate radial density profiles even after adding sufficiently noise to our synthetic radiographs to make inverse Abel transforms untenable. Compared to state-of-the-art classical QUBO solvers, GUROBI and the Hamze-Freitas-Selby algorithm, the quantum D-Wave 2X was orders of magnitude faster for the same final accuracy. Our results indicate a potential strategy for integrating neuromorphic and quantum computing techniques.**

***Index Terms*—Sparse coding, Lateral inhibition, Compressive sensing, Radiographs, NP-hard, Quantum D-Wave 2X machine, Quantum annealing**

I. Introduction

Inferring 3D depth from 2D retinal images is a fundamental task of the visual cortex. Computer simulations suggest that unsupervised training of a model of V1 simple cells in an manner designed to optimize the sparse reconstruction of 2D stereo images or video naturally supports the acquisition of 3D depth-selective responses [1], [2]. In these studies, models of V1 simple cross were implemented using a locally competitive algorithm (LCA), a neurally plausible mechanism for inferring sparse representations based on lateral inhibition between leaky-integrator neurons governed by a soft-threshold transfer function [3]. Based on their somewhat analogous architectures, we postulated that a version of LCA could be efficiently implemented on the D-Wave 2X quantum annealing computer. As a more tractable problem suited for the limited size of the D-Wave 2X, we considered the inference of 3D density profiles from synthetically generated 2D radiographs of mathematically defined objects. When sufficient noise is added to the synthetic radiographic images, density reconstructions based on direct inverse transform methods become wildly inaccurate. In contrast, we found that implementing a model of binary V1 simple cells on the D-Wave 2X allowed us to infer accurate radial density profiles even in the presence of similar noise levels. Our results suggest a general contractor strategy for synthesizing neuromorphic and quantum techniques.

II. Neuromorphic approach

A. Sparse representations

The hypothesis that neurons encode stimuli by inferring sparse representations explains many of the response properties of simple cells in the mammalian primary visual cortex [4], [5]. Given an overcomplete, non-orthonormal basis $\{\phi_i\}$, inferring a sparse representation involves finding the minimal set of non-zero activation coefficients $\boldsymbol{a}$ that accurately reconstruct a given input signal $\boldsymbol{I}$, corresponding to a minimum of the following energy function:

$$E(\boldsymbol{I}, \boldsymbol{\phi}, \boldsymbol{a}) = \min_{\{\boldsymbol{a}\}}[\ \frac{1}{2}||\boldsymbol{I} - \boldsymbol{\phi a}||^2 + \lambda||\boldsymbol{a}||_0\ ] \qquad (1)$$

where λ is a trade-off parameter that determines the balance between reconstruction error and the number of non-zero activation coefficients. A larger λ serving as a higher threshold will result in a more sparse representation solution to problem Eq. (1). This energy function is non-convex and contains multiple local minima, so that finding a sparse representation falls into an NP-hard complexity class of decision problems [6].

B. Lateral inhibition

Previous work [3] has shown that *sparse coding* optimization problems can be solved using the dynamics of neural networks incorporating lateral inhibition, a biologically

plausible implementation of a sparse solver referred to as an locally competitive algorithm, LCA. LCA can be summarized by the following elements. The output or firing rate of a neuron is given by a threshold or activation function of its membrane potential. The output of each neuron (usually zero) corresponds to the activation coefficient of its associated basis vector. The input to each neuron is given the inner product of its feature vector with the stimulus. The magnitude of the lateral inhibition between any pair of neurons is determined by the inner product of their corresponding basis vectors. *A sparse solution is found by allowing the resulting dynamical system to evolve to a minimum energy configuration*. However, such dynamical systems are susceptible local minima. Here, we use quantum annealing to explore the global energy landscape. In qualitative terms, we associate the input to each binary all-or-none neuron to the bias or external field applied to each spin or qubit and the lateral inhibition between neurons with the spin-spin coupling between superconducting qubits inside the quantum D-Wave machine hardware [7]. An optimal solution to the sparse coding problem is then determined by the configuration of spin orientations that minimize the total energy to the associated Ising system.

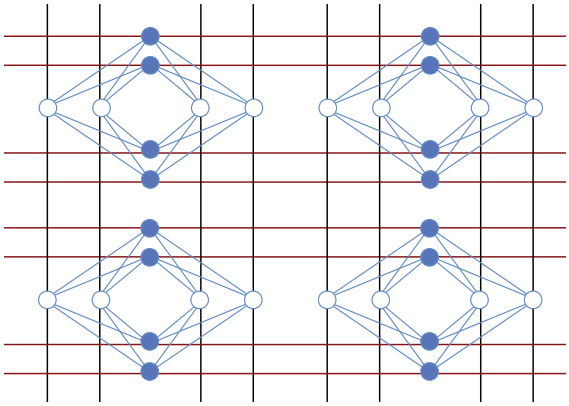

Fig. 1: (Color online) A subset of the Chimera D-Wave consisting of 32 qubits arranged in the form of 2x2x8 in which 8 qubits in each of the 4 unit cells plotted here 1) are grouped into two arrangements of vertical (horizontal) orientations as blue (white) circles and 2) interact with one another through the 16 bipartite interactions (blue edges). Nearest neighboring bipartite interactions between each pair of nearest neighboring unit cells are characterized as black and red edges.

III. Mapping sparse coding onto QUBO D-Wawe 2X

A. Lateral inhibition neurons versus quantum bits: transformation relations

We investigate the application of quantum annealing to the solution of sparse coding quadratic unconstrained binary optimization (QUBO) problems employing binary activation coefficients that arise in the context of sparse coding. Each neuron is mapped to a qubit in the Chimera hardware and because a qubit is assigned only two possibilities 0 and 1 (or equivalently in an Ising problem, -1 and 1 for spin orientation) the neuron in our case is treated as a "quantum object" bearing two possibilities: firing with maximum activation with coefficient 1 and silent with coefficient 0. Because each neuron is a quantum object, the state of any neuron is described in general by a superposition of 1 and 0, in which the neuron is both active and non-active at the same time, a logical impossibility for any classical system. It is this quantum superposition that allows the D-Wave to explore the entire energy landscape at once.

The D-Wave 2X [7] finds optimal solutions to a (discrete) Ising system consisting of N_q binary variables via quantum annealing. Such N_q-body systems can be described by the following classical Hamiltonian:

$$H(\boldsymbol{h},\boldsymbol{Q},\boldsymbol{a}) = \sum_{i}^{N_q} h_i a_i + \sum_{i<j}^{N_q} Q_{ij} a_i a_j \qquad (2)$$

with binary activation coefficients $a_i = \{0,1\}\,\forall i \in (1,2,3,...,N_q)$. This objective function defines a QUBO problem. We cast our sparse coding problem, Eq. (1), into QUBO form, Eq. (2), by the transformations [8]:

$$\begin{aligned} h_i &= (-\boldsymbol{\phi}^T \boldsymbol{I} + (\lambda + \frac{1}{2}))_i, \\ Q_{ij} &= (\boldsymbol{\phi}^T \boldsymbol{\phi})_{ij}. \end{aligned} \qquad (3)$$

In Eq. (3), the bias term $\boldsymbol{h}$ (elements h_i) in the Ising model is proportional to the weighted input $\boldsymbol{\phi}^T \boldsymbol{I}$ while the coupling term $\boldsymbol{Q}$ (elements Q_{ij}) corresponds to lateral competition (see also [3]) between qubits given by the interaction matrix $\boldsymbol{\phi}^T \boldsymbol{\phi}$ [with self-interaction excluded and $\boldsymbol{Q}$ being symmetric i.e. $Q_{ij} = Q_{ji}\ \forall i \neq j$, see Eq. (2)].

B. D-Wave 2X hardware

The D-Wave 2X [7] consists of 1152 qubits arranged into 12x12 unit cells, forming a Chimera structure with dimensions 12x12x8. Sparse interactions between qubits are restricted to the 16 connections within a unit cell and the 16 connections between nearest-neighboring unit cells [7].

In detail, each unit cell contains 4 qubits aligned along a horizontal axis and 4 qubits aligned vertically (see Fig. 1). Within a unit cell, the 4 qubits of a given orientation can only can interact with the 4 qubits with the opposite orientation (e.g. see the two groups of white and blue filled circles represented as qubits with different geometries in Fig. 1). Between unit cells, interactions are only allowed between nearest-neighbors (blue edges) and even between nearest-neighbors the allowed connections are restricted according to relative orientation. A vertically (horizontally) oriented qubit can only connect to the two vertically (horizontally) oriented qubits at the same relative position in the nearest-neighboring unit cells immediately above (left) and below (right) along one column

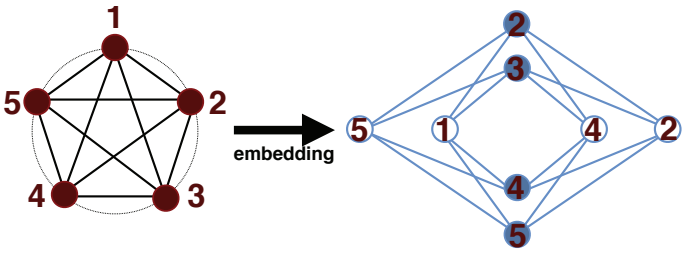

Fig. 2: (Color online) An illustrating example for mapping a 5 fully-connected particles (left, 5 red filled circles with 10 black edges) onto a part of the Chimera hardware using 8 qubits (right, 8 blue circles) with all eligible connections (blue edges).

of the 12x12 grid [see black (red) edges in Fig. 1]. Thereby, in a chimera graph, one qubit can interact with at most 6 other qubits.

Embedding technique

Despite the sparsity of physical connections on the D-Wave, it is nonetheless possible to construct graphs with arbitrarily dense connectivity by employing "embedding" techniques. Embedding works by chaining together physical qubits so as to extend the effective connectivity but at the cost of reducing the total number of available logical qubits. Because logical qubits do not need to follow the connection rules that physical qubits do, it is possible to implement general QUBO problems with arbitrarily dense connectivity on the sparsely connected D-Wave chimera. The D-Wave API provides a heuristic algorithm that searches for an optimal embedding that minimizes the number of physical qubits that are chained together.

For illustration purposes, we plotted in Fig. 2 a cartoon describing a simple embedding example mapping a fully-connected problem containing $N_q = 5$ "logical" qubits (red circles) interacting via 10 bipartite couplings (10 black edges) to a Chimera graphed problem. The corresponding embedded output solution, shown on the right-hand-side of Fig. 2, is a subgraph consisting of 8 "physical" qubits (blue circles), numbered from 1 to 5 denoting the corresponding 5 "logical" qubits. Some physical qubits are tied together by strong negative couplings to create and effective logical qubit with greater connectivity. Note that this embedding is not unique. As the number of constructing particles N_q increases ($20 < N_q \leq 47$ as studied in the remaining part of the paper), each "logical" qubit now can interact with a much denser number of neighboring qubits, which results in a much more sophisticated embedding solution.

The exact mapping of a spin glass problem onto the physical D-Wave 2X chimera, including defects, can typically contain approximately $N_q \sim 1000$ spins (qubits) with > 3000 local spin-spin interactions. In contrast, embedding an arbitrary QUBO problem onto the same 2X chimera typically allows no more than $N_q \sim 47$ nodes (logical qubits) but these nodes may be fully connected. Thus, embedding effectively trades qubits for connectivity.

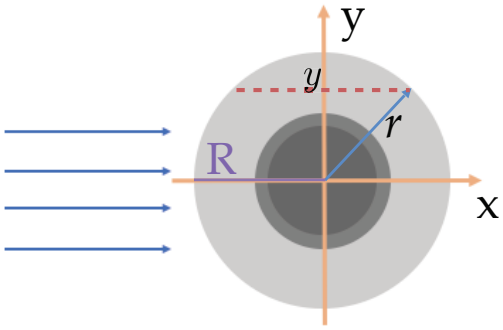

Fig. 3: (Color online) A cross-section cartoon describing how the Abel integral is formed. A set of parallel beams (blue arrows) passes through an axially-symmetric object, forming a radiographic projection $\boldsymbol{I}$ that varies along the y-axis.

IV. Radiographic Inference on the D-Wave

A. Compressive sensing formalism

Compressive sensing [9], [10] seeks to accurately reconstruct a compressible signal $\boldsymbol{f}$ of dimension Nx1 using a (much) smaller number of measurements M compared to the dimensionality of the original signal ($M \leq N$). The assumption of compressibility can be formulated as:

$$\boldsymbol{I} = \boldsymbol{A}\,\boldsymbol{f}. \tag{4}$$

Here, $\boldsymbol{A}$, of dimension MxN, is called the sensing matrix for measurement $\boldsymbol{I}$ of dimension Mx1. The reconstruction $\boldsymbol{I}$ of signal $\boldsymbol{f}$ is accurate when $\boldsymbol{f}$ is known to be sparse in some basis $\boldsymbol{F}$ such that $\boldsymbol{f} = \boldsymbol{F}\boldsymbol{a}$ then we can rewrite:

$$\boldsymbol{I} = \boldsymbol{\phi}\boldsymbol{a} \tag{5}$$

with

$$\boldsymbol{\phi} = \boldsymbol{A}\,\boldsymbol{F}. \tag{6}$$

Thus, to infer a radial density profile, we seek a sparse representation $\boldsymbol{a}$ for signal $\boldsymbol{f}$ that recovers the measurement $\boldsymbol{I}$. In particular, we minimize an energy function of the sparse code form with sparse basis $\boldsymbol{\phi}$. If the activation coefficients $\boldsymbol{a}$ are assumed to be binary, we have cast radiographic inference as a QUBO problem that can be solved on a quantum D-Wave 2X machine using embedding techniques. A more complete discussion on compressive sensing is specialized in [9], [10].

B. Abel transformation

Assume that an object possesses axial symmetry (e.g. a cylinder with radius R) and a radial density dependence given by $\boldsymbol{f} \equiv f(\overrightarrow{r})$ with $\overrightarrow{r}$ denoting radial position. The Abel transform [11] yields the log of a radiographic projection where a set of parallel beams (blue arrows in Fig. 3) passes through an axially-symmetric object forming a radiographic

image projection $\boldsymbol{I}$ that varies along the y-axis. Formally, we write:

$$\boldsymbol{I} = \boldsymbol{A}\boldsymbol{f} \tag{7}$$

where operator $\boldsymbol{A}$ now represents an Abel transformation. For implementation on the D-Wave, we cast $\boldsymbol{A}$ as a matrix of dimension MxN obtained by discretizing the Abel transform as:

$$A_{ij} = 2\sum_j \frac{r_j \Delta_{ij}}{\sqrt{r_j^2 - y_i^2}} \tag{8}$$

where $i = 1, 2, ..., M$, $j = 1, 2, ..., N$, Δ_{ij} is the discretized integration step size and the upper limit corresponds to $y_M = r_N = R$.

C. Input for D-Wave 2X

We still need to identify a basis in which the radial density function $\boldsymbol{f}$ is sparse. Here, we assume that the synthetic radial density profile $\boldsymbol{f}$ is sparse in a Fourier domain. In general, $f(\overrightarrow{r})$ can be expanded as a discrete Fourier series F_n:

$$f(r_j) = \sum_n^{N_q} a_n F_n(r_j) \tag{9}$$

where the $F_n(r_j)$ are discrete cosine functions forming Fourier matrix $\boldsymbol{F}$ of dimension NxN_q. The Fourier coefficient vector $\boldsymbol{a}$ is synthesized such that 1) $\boldsymbol{a}$ is sparse; most of the Fourier coefficients a_n in $\boldsymbol{a}$ are zero and 2) $\boldsymbol{f} = \boldsymbol{F}\boldsymbol{a}$.

We now can now rewrite

$$\boldsymbol{I} = \boldsymbol{A}\boldsymbol{F}\boldsymbol{a} = \boldsymbol{\phi}\boldsymbol{a}, \tag{10}$$

with $\boldsymbol{\phi}$ now defined as

$$\boldsymbol{\phi} = \boldsymbol{A}\boldsymbol{F}, \tag{11}$$

in the form of a sparse coding compressive sensing representation [9], [10]. Note that $\boldsymbol{\phi}$, in this case defined via Eqs. (8) and (9), which is equivalent to the representation $\boldsymbol{\phi}$ in Eq. (6), can be interpreted as "dictionary" for a sparse coding implementation. The sparse set of Fourier coefficients $\boldsymbol{a}$ can then be determined by minimizing the objective function $||\boldsymbol{I} - \boldsymbol{\phi}\boldsymbol{a}||^2 + \lambda||\boldsymbol{a}||_0$ [defined as in Eq. (1)]. The size of $\boldsymbol{a}$ is limited by the capacity of the D-Wave 2X, which in the case of embedding a fully connected graph is $N_q = 47$. In the current context, the Hamiltonian parameters $\boldsymbol{h}$ and $\boldsymbol{Q}$ defined above in Eq. (3), respectively, become: $\boldsymbol{h} \sim -\boldsymbol{\phi}^T\boldsymbol{I} \sim -(\boldsymbol{A}\boldsymbol{F})^T\boldsymbol{I}$ and $\boldsymbol{Q} \sim \boldsymbol{\phi}^T\boldsymbol{\phi} \sim (\boldsymbol{A}\boldsymbol{F})^T(\boldsymbol{A}\boldsymbol{F})$.

V. Evaluating performance of the quantum D-Wave 2X versus the state-of-the-art classical solver GUROBI

We use the following solvers/algorithms to find solutions to our QUBO problem: the D-Wave 2X, GUROBI, an exact classical solver, and Abel inversion, which applies to continuous variables. During the cooling process, approximately 4.5% (for the current D-Wave 2X hardware we used) of the physical qubits become unusable. The D-Wave 2X at Los Alamos National Laboratory on which these runs were performed had 1100 active qubits with 3068 active couplers. Fig. 4a shows the radial density profile obtained from the quantum D-Wave 2X machine. We used $\lambda = 23$ for $N_q = 47$ nodes. The synthesized density profile is relatively well reconstructed and all non-zero Fourier coefficients were retrieved.

Adding noise to the synthetic radiograph did not affect the density reconstruction (red and dashed lines coincided) and all non-zero Fourier coefficients were again retrieved. In contrast, the density profiles produced by the inverse Abel transform, although exact in the noiseless case (not shown), were severely impacted by the addition of noise (Fig. 4b).

GUROBI [12], a commercial optimization tool, yields exact solutions to general QUBO problems. The D-Wave 2X was able to find all non-zero Fourier coefficients within a minute, including the time for all processing steps, such as i/o, embedding, annealing, reading, etc. When applied to same radiographic-inspired QUBO using e.g. $N_q = 47$ nodes, however, GUROBI required several hours to converge and in some cases was unable to obtain the exact solution before a 10 hour cutoff was imposed.

We have so far examined the cases $N_q = 47$, the maximum number of logical qubits we could embed on the quantum D-Wave 2X hardware. Finding a global minimum to our densely connected QUBO problem should become more difficult as N_q becomes larger, since the number of possible combinations of Fourier coefficients that can reproduce the synthetic radiograph grows combinatorially. In this section, we study a range of $N_q = 21, 25, 30, ..., 37$ (Fig. 5) with λ slightly increasing as N_q increases to maintain a similar sparsity scaled to N_q. Below approximately $N_q = 40$, GUROBI is able to find an optimal solution within a tractable amount of time. However, above approximately $N_q = 40$, GUROBI can no longer find a global minimum within the 10 hour cutoff imposed for this study. For values of N_q ranging between $N_q = 21$ to $N_q = 37$, runtime grew exponentially as a function of N_q. A linear fit (red dashed line in Fig. 5) yielded a slope of ~ 0.27, which is close to $\log_{10} 2$, as expected for an exact but exhaustive combinatorial search.

VI. Other comparisons

A. Abel inversion

Abel inversion has been known to recover signals in the absence of noise. We observed that Abel inversion produced an absolute squared error $\Delta f = ||\boldsymbol{f} - \boldsymbol{F}\boldsymbol{a}||^2 \ll 0.1$. The absolute error rate for the D-Wave 2X is around 0.54 [see Fig. 4(b)]. Thus, Abel inversion is much more precise than the D-Wave 2X in the absence of noise. Note that the quantum D-Wave solver is binary while standard Abel inversion gives an exact continuous solution. However, when approximately 20% white noise is added to both the radiographic projection $\boldsymbol{I}$ and to the original synthesized radial density profile $\boldsymbol{f}$, the inverse Abel transform performs poorly. While the Abel inversion is very sensitive to noise, we found that the D-Wave is robust to noise below about $< 25\%$ of the signal.

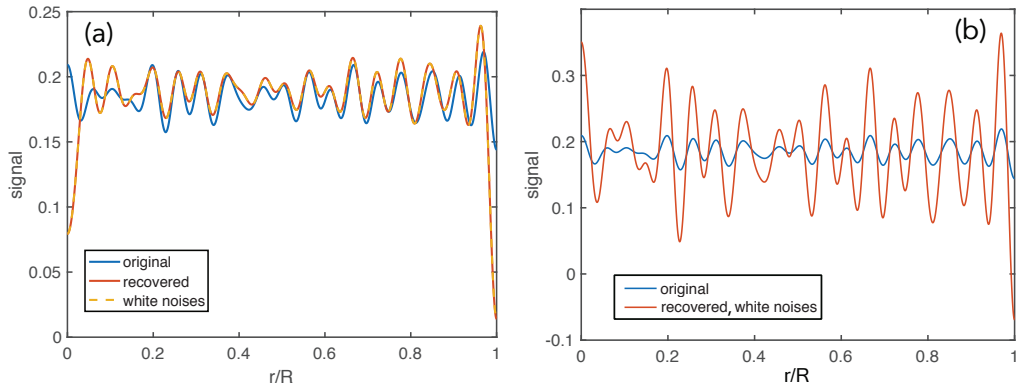

Fig. 4: (Color online) (a): Radial density solutions. Solid lines refer to the true radial density (blue) and D-Wave reconstructed radial density (red). Dashed line refers to the reconstructed radial density in the presence of noise. (b): Radial density obtained using Abel inversion. Blue line refers to the original radial density and red the reconstructed radial density in the presence of noise.

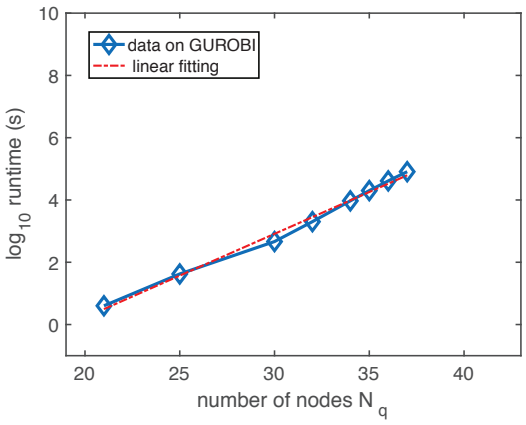

Fig. 5: (Color online) Time for GUROBI to obtain an exact solution of a densely connected QUBO problem as a function of N_q, the total number of possible Fourier coefficients. Time axis in log base 10.

B. Performance with the Chimera-inspired solver Hamze-Freitas-Selby (HFS)

Lastly, we compare the D-Wave 2X to the classical Hamze-Freitas-Selby algorithm (HFS) [13], [14]. HFS is a QUBO solver implemented by Selby [15] that uses the D-Wave Chimera structure to successfully solve spin-glass problems. In HFS, low-bandwidth graphs are used instead of individual spin-flips to search for optimum solutions among all possible configurations.

Previous work has shown that HFS performs well compared to the D-Wave on spin-glass problems when the number of qubits is limited [15]. We used the concept of *time to solution*, the main runtime measure for HFS, which in a D-Wave machine is defined as the ratio of *annealing time* and *probability to be the optimum solution*. Annealing time for the quantum D-Wave 2X hardware we used is $\sim 20\mu s$. We report that for this problem type, *time to solution* in the D-Wave 2X is a few hundreds times faster on average than that for HFS. The *time to solution* of one of our inputs using the best combination of HFS parameters we found (we tried many different combinations) for this case was about ~ 0.8 seconds (equivalently $\sim 0.8\,10^6\mu s$) while *time to solution* was around $< 10^3\mu s$ on D-Wave 2X. Note that this *time-to-solution* measure is not the wall-clock time as was used for the comparisons with GUROBI in the previous section. For example, in this case, it took about 6 minutes for HFS to yield its optimum solution according to the wall-clock time while it was about less than 1 second for D-Wave 2X, respectively. We performed our work on Intel(R) Xeon(R) CPU E5-2667 v2 @ 3.30GHz.

VII. Conclusion

Radiographic inference is an important technique for remotely interrogating both static and dynamic physical objects. Because a single radiograph can only yield a single cross-section, inferring the corresponding 3D structure is in general a difficult, ill-posed inference problem. However, this inference problem can be rendered more tractable to the extent that prior information about the physical objects being imaged can be brought to bear. For example, an Inverse Abel transform utilizes the prior information that the underlying objects possess cylindrical symmetry. Here, we incorporate prior information using a sparse coding framework, which involves finding a basis in which the 3D structure of the underlying objects can be recovered in terms of a small number of basis coefficients. In general, if a sufficient number of training examples spanning the range of possible physical realizations are available, such a basis can always be found using standard dictionary learning techniques. Here, as an expedient, we defined radially

symmetric objects that were constructed to be sparse in a Fourier domain. However, the underlying principle should extend to the interpretation of radiographs of any set of objects that can be sparsely represented in some basis.

We used a biologically plausible approach to sparse coding based on lateral inhibition to infer radial densities from radiographic projections. By employing binary neurons, we mapped the radiographic inference problem to a QUBO problem, a format that allowed direct implementation on the D-Wave quantum annealing computer. We used quantum annealing to explore the global energy landscape of a 47-neuron systems governed by lateral competition to search for a minimal sparse representation of firing neurons that best reconstruct our synthesized radial density profile input. We compared the performance of the quantum annealing D-Wave 2X computer on a densely connected QUBO problem against several classical approaches (GUROBI, Abel inversion, HFS). Previous comparisons of D-Wave quantum annealing computers against classical algorithms focused on spin-glass problems mapped directly onto the sparsely connected D-Wave chimera. We found that the D-Wave 2X hardware dramatically outperforms GUROBI for these densely connected QUBO problems, obtaining solutions within one minute on problems that GUROBI was unable to solve in under the 10 hour cutoff imposed for this study. As a function of the number of binary coefficients N_q from which the solver much chose, the D-Wave 2X exhibited clear superiority over GUROBI for approximately $N_q \geq 37$. We also made a comparison for our results with a Chimera-inspired solver HFS. Again, the D-Wave 2X outperformed HFS considerably, being approximately 100 times faster than HFS for this problem. We showed that the D-Wave 2X is able to filter white noise added to the synthetic radiograph whereas the same added noise levels causes Abel inversion to perform poorly.

The availability of the D-Wave 2X [7] quantum annealing computer naturally encourages comparisons with classical approaches. For example, Vinci *et. al.* [16] compared D-Wave performance on spin-glass problems against the classical Hamze-Freitas-Selby algorithm [13], [14] and showed that quantum annealing is $2-3$ orders of magnitude faster. However, by targeting spin glass problems that map directly onto the sparsely connected D-Wave chimera, such comparisons do not address the relative performance of quantum annealing for the solution of general QUBO problems, particularly those problems which involve dense connectivity. Here, we start with a general QUBO problem inspired by a compressive sensing [9], [10] approach to radiographic inference and use embedding tools to map this problem onto the D-Wave Chimera

In conclusion, we find that quantum annealing provides a strategy for inferring the sparse set of binary coefficients used in the object definition by finding good local minima of a non-convex cost function which penalizes the differences between the actual and reconstructed radiographs. Similar to what has been reported for spin glass problems, the D-Wave 2X is found to exhibit orders of magnitude better performance even on a QUBO problem with dense connectivity that does not match the D-Wave chimera. The D-Wave 2X yielded the lowest energy solutions in a few seconds with all the synthesized non-zero binary coefficients recovered. The D-Wave solution remained robust after white noise was added to the radiograph, a manipulation that causes the closed-form inverse Abel transform to perform poorly. Our results suggest an important application domain based on sparse coding and compressive sensing in which quantum annealing can rapidly find solutions to NP-hard problems that pose considerable difficulties for existing state-of-the-art solvers based on classical algorithms.

ACKNOWLEDGMENT

N.T.T.N greatly acknowledges Carleton Coffrin for providing technical assistances regarding GUROBI optimization and for helpful discussion. We thank Pete Schultz and John Perry for fruitful discussions.

REFERENCES

[1] S. Y. Lundquist, D. M. Paiton, P. Schultz, G. T. Kenyon, "Sparse Encoding of Binocular Images for Depth Inference," in Proceedings of the Southwest Symposium on Image Analysis and Interpretation (SSIAI) 2016, pp. 121–124, March 2016.

[2] Sheng Y. Lundquist, Melania Mitchell, and Garrett T. Kenyon, "Sparse Coding on Stereo Video for Object Detection," arXiv:1705.07144, May 2017, in Proceedings of the Workshop on Learning with Limited Labeled Data: Weak Supervision and Beyond, NIPS, December 2017.

[3] C. J. Rozell, D. H. Johnson, R. G. Baraniuk, and B. A. Olshausen, "Sparse Coding via Thresholding and Local Competition in Neural Circuits," Neural Computation, vol. 20, pp. 2526–2563, October 2008.

[4] B. A. Olshausen and D. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," Nature, vol. 381, pp 607–609, June 1996.

[5] B. A. Olshausen and D. Field, "Sparse Coding with an Overcomplete Basis Set: A Strategy Employed by V1?" Vision Res., vol. 37, pp. 3311–3325, December 1997.

[6] B. Natarajan, "Sparse approximate solutions to linear systems," SIAM Journal on Computing, vol. 24, pp. 227–234, April 1995.

[7] D-Wave systems: http://www.dwavesys.com/

[8] N. T. T. Nguyen and G. T. Kenyon, "Solving sparse representation for object classification using quantum D-Wave 2X machine," in Proceedings of The First International Workshop on Post Moore's Era Supercomputing, pp. 43–44, November 2016.

[9] E. J. Candès, "An Introduction to Compressive Sensing," IEEE Signal Processing Magazine, vol. 25, pp. 21–30, March 2008.

[10] D. Donoho, "Compressed sensing," IEEE Trans. Inform. Theory, vol. 52, pp. 1289–1306, April 2006.

[11] R. N. Bracewell, "Numerical transforms," Science, vol. 248, pp. 679–704, May 1990.

[12] http://www.gurobi.com

[13] F. Hamze and N. de Freitas, "From fields to trees," in Proceedings of the 20th Conference in Uncertainty in Artificial Intelligence, Banff, Canada, pp. 243–250, July 2004.

[14] A. Selby, "Efficient subgraph-based sampling of Ising-type models with frustration," arXiv:1409.3934, September 2014.

[15] A. Selby, https://github.com/alex1770/QUBO-Chimera

[16] W. Vinci and D. Lidar, "Optimally Stopped Optimization," Phys. Rev. Applied, vol. 6, pp. 054016-1–054016-22, November 2016.